ACCELERATING INTELLIGENCE OF AUTONOMOUS ROBOTIC SPACECRAFT

MICRO-14 Session Summary

Background and Introduction

A team of motivated stakeholders and subject-matter experts from government, industry, and academia organized and executed an interactive workshop addressing "Accelerating Intelligence of Autonomous Robotic Spacecraft" — this was a follow-on to a panel session held at 2021 ASCEND on "Advancing Robotic Space Exploration."The team identified a set of four specific topics and a set of motivating questions for each topic to structure the conversation at the workshop.

This summary includes a brief description of the theme, an overview of the topic discussions, and a summary of proposed next steps for the AIAA spacecraft autonomy community. These candidate accelerating actions are intended to be actionable steps that will provide the basis for continued development and evolution of autonomous robotic spacecraft.

Description

There is massive worldwide investment in artificial intelligence and autonomous systems across a variety of commercial and governmental domains. How can this interest and investment be best leveraged to advance the capabilities of autonomous robotic spacecraft? This workshop explored approaches to building a public–private partnership to accelerate deployment of these capabilities. The participants examined what products or results would be of most benefit to accelerate capabilities of autonomous spacecraft, and of relevant Earth analogs (e.g., underwater exploration)?

Topics

1. Autonomy in Mission Formulation
2. Infusion of Autonomy Capability
3. Digital Engineering for Autonomy Verification and Validation (V&V)
4. Explainability and Human-Machine Teaming

TOPIC 1: AUTONOMY IN MISSION FORMULATION

This topic area addresses 1) the new mission sets enabled by the availability of greater autonomy, and how mission formulation practices may/should be adjusted to take advantage of emerging AI and autonomy capabilities, 2) the importance of multi-mission autonomy architectures and challenges to their broad deployment, and 3) the role of autonomy capabilities in mitigating operational risks of missions.

DISCUSSION QUESTIONS:

Autonomy Mission Concepts

What new mission sets are enabled by the availability of greater autonomy? How do we break the cycle of overly relying on past experience to inform future mission concepts and architectures? How should mission formulation organizations and processes change to take full advantage of the promise of AI and autonomy? New space system designs incorporating autonomy may need to be appropriately bold or at least different from historical designs, as mission objectives continue to evolve. How do we know when a new space mission concept is different, in features and/or degree, and that the concept is relevant and feasible?

Autonomy Core Architectures

Most practitioners agree there are core autonomy capabilities that appear again and again across space mission designs. Assuming such capabilities could be captured and maintained as rigorous, stable product lines, are there ways to make them available broadly to the spacefaring community, balancing openness with competitiveness?

Autonomy for Ongoing Validation

Autonomy may be proposed for space missions in part to address the operational risk of imperfectly known environments. How may autonomy capability itself be leveraged to address operational risks as they are disclosed or clarified during operations, by selecting, adjusting, or evolving system behaviors as the mission unfolds?

TAKEAWAYS:

Autonomy Mission Concepts

- It is natural in space mission design to build on what has been already successful, adapting to new mission objectives and requirements. This approach is a tried-and-true practice for managing risk. But this approach begs the question of whether mission concepts centrally incorporating autonomy would/should be different from traditional mission concepts.
- To address this question, two elements are needed: 1) a way of measuring how different one mission concept is from another, and 2) a way of evaluating mission concepts for success.
    - o An evaluative framework for mission concepts should include criteria such as cost and risk, as well as criteria related to functional and performance success, dependent on the degree of onboard autonomy used in the mission.
    - o Machine learning (ML) techniques, e.g., classification, may help determine differences among mission concepts—mapping concepts to a feature-based metric mission design space.
- At the moment, we do not know if mission concepts tend to look similar because they actually do represent best solutions, or whether we may be missing different, perhaps better mission concepts because we are remaining, unwittingly or not, in a comfort zone.
- Artificial intelligence techniques, seen as not just contributing but as defining aspects of space autonomy, can be considered both for their contributions to the functionality of the resulting mission systems, and to the methodologies for developing those systems—both across a given mission lifecycle, and across multiple missions. This potential for wide impact of AI becomes evident when considering autonomy in the context of mission formulation.

Further Questions

- What set of features are needed to describe the design space of mission concepts incorporating autonomy?

- What process or expertise updates may be needed to properly evaluate mission concepts incorporating autonomy in competitive solicitations?
- What should be the role of modeling and simulation in the evaluation of mission concepts?

Autonomy Core Architectures

- The core functions of a space mission are to a degree dependent on factors such as operating regime and payload. A mapping or surveillance satellite in Earth orbit will have different functionality needs than a deep space telescope at L2, although both depend critically on GN&C capability.
- Parameterized design solutions may offer a way to address how to capture core autonomy functions in software product lines while maintaining flexibility for differentiated mission functionality needs. Algorithms and software modules should be designed with such generality in mind.

Further Questions

- How to share core autonomy architectures, balancing openness with competitiveness?
- How to allocate costs for generalized (product line) solutions, beyond the requirements of a specific mission?

Autonomy for Ongoing Validation

- A deep, sometimes unspoken motivation for incorporating autonomy in a mission concept is to be able to grapple with operational uncertainty. Mars rovers provide good examples, as operations entail moment-to-moment interaction with the environment in the presence of light-time delay.
- An interesting question is whether AI/autonomy itself might provide solutions. An onboard planner powered by more capable flight computing could precompute contingency plans, invoking them as needed as operational realities intrude, rather than halting and dropping into another round of planning. Such a planner might access or generate a safe behavior envelope for the platform, consistent with the sensed environment.
- One can view such an approach as continuous validation (of proposed actions in context). In this way, validation extends to operations, beyond launch and deployment. The autonomous platform contributes to ongoing operational risk assessment and management.

Further Questions

- What guarantees of correct and/or safe behavior are possible for autonomous systems?
- In addition to advances in flight computing, what other adjacent technology advances can be enabling for autonomy?

CANDIDATE ACCELERATING ACTIONS:

Autonomy Mission Concepts

- Investigate ML (classification) as a means to determine distances among mission concepts.
- Update evaluative frameworks for mission concepts incorporating autonomy, extending to simulated functional performance tests.

Autonomy Core Architectures

- Engage an appropriate program office to fund autonomy software product line solutions.
- Form a task force to generate guidelines for open strategies to develop, verify, share, and maintain autonomy software product lines.

Autonomy for Ongoing Validation

- Investigate how research in trusted autonomy and human-machine teaming should be leveraged in support of autonomy for ongoing validation.
- Engage a conversation with stakeholders on how autonomy can support risk management.

TOPIC 2: INFUSION OF AUTONOMY CAPABILITY

This topic area addresses the key challenges to infusing autonomy capability into space systems.

DISCUSSION QUESTIONS:

What are some of the fundamental challenges in deploying increasingly autonomous systems? Are current engineering processes and capabilities up to the challenge?

A common issue brought up as a challenge to adoption of autonomy is verification and validation. What techniques can be used to increase our ability to certify autonomous systems? Which stakeholders will be the most difficult to convince?

What is the appropriate role of early adopters?

What kinds of opportunities are surfacing for demonstrating and maturing autonomy capabilities? What is the role of on-orbit testbeds?

TAKEAWAYS:

The developed issues fell into two broad themes. One theme was issues related to establishing trust that autonomous systems would perform as designed. The other theme was how the various unique limitations of the space environment and/or space systems cause issues for autonomous systems. Around trust, the key issues include:

- Verification: not all possible scenarios or outcomes can be tested
- Security, including cyber concerns and being able to scale security measures
- Limited training data available for machine learning systems
- Low artificial intelligence maturity
- User resistance to adoption

Limitations unique to space systems include:

- Hardware capabilities, particularly computing power
- Communication limitations, particularly a lack of connection to the internet and cloud services
- Lack of software modularity (unique space software with little reuse)

CANDIDATE ACCELERATING ACTIONS:

Trust:

The standardization of verification expectations, test cases, and test beds would build confidence in the verification process. Interoperable environments would be a benefit, so that autonomous tools could be observed working well together. The use of fast prototyping methods would provide for more test data. The use of already verified systems from another domain (potentially non-aerospace) would build confidence. Care is needed to evaluate input data for machine learning to ensure the data is relevant.

Limitations unique to space systems:

Possible solutions for the limitations of space hardware and environments included increasing and using both funding and mass margin to design more robust systems. The selection of space applications that are well suited to autonomy is also critical. Using more standardized hardware and data, and designing architectures to accommodate commercial-off-the-shelf technology, will build up capabilities for autonomous systems. Other ideas included leveraging digital twins on Earth to analyze behaviors, and the use of cloud computing in the space environment to provide sufficient computing resources for autonomy.

TOPIC 3: DIGITAL ENGINEERING FOR AUTONOMY V&V

This topic addresses the V&V of autonomous systems, and the role of digital engineering techniques in providing solutions.

DISCUSSION QUESTIONS:

How can the convergence of design of experiments (DoE) and live/virtual/constructive modeling and simulation advance V&V for autonomous space systems?

What are the best communities to leverage DoE expertise from?

Are there APIs or standards, e.g., from the modeling/simulation area, that could help (or hurt)?

Shared data, tools, and analytic results are key. How can AIAA support this (workshops, conferences, special sessions, etc.) during ASCEND?

TAKEAWAYS:

The team discussing this issue developed two perspectives to help guide the identification of solution areas:

1. Autonomy and digital engineering are not "new" ideas, but using them in concert for space applications comes with challenges that need to be overcome, including:
   - Limited data available for checking model correctness.
   - Challenges from "launch and never touch again"; current "source of truth" unavailable for inspection.
   - Lack of transparency (particularly for data-driven approaches like machine learning).
   - Lack of comprehensive, standardized tool kits for digital engineering of autonomous systems.
   - Competition for these skill sets from many other industries. How do we build these skill sets in aerospace schools? How do we do in-house training efficiently and effectively for our personnel?
2. The space community needs to work on software and test engineering:
   - Leverage modularity, reusability, and other software "-ilities" for both testing and on-orbit autonomy software.
   - Better understand the sensitivities of models used for the space domain to "real vs synthetic" data, and design robustness into the models and autonomy capabilities.
   - Software testing as a whole in the space community is lacking.

CANDIDATE ACCELERATING ACTIONS:

- Engage with other autonomy areas and related fields: UAS, automobile, healthcare (clinical trials), etc. Learn from these other communities how trust is developed in safety-critical systems. Bring more "space outsiders" into the community and cross-train them to leverage their different perspectives in tackling our problems.
- Identify a lead for autonomy testing standards in general and specifically for space across DoD, IEEE, CCSDS, AIAA, etc., and work toward a consortium for building needed standards. These standards would include a definition of "spaceworthiness," similar to the airworthiness criteria for aircraft.
- Define best practices for space software development, sustainment, and testing, especially for autonomy software. These best practices should leverage lessons from other communities who have experience in related issues. Define a basis for comprehensive understanding of autonomous system behavior, based on analyses and testing.
- Explore options to develop an anonymized autonomy lessons learned body of knowledge, which includes best practices and what not to do, without specifics on individual failures and other proprietary information.
- Professional societies should define autonomy challenge problems and facilitate open-source collaboration.
- Work with young professionals, academic partners and enthusiast groups to develop a smallsat autonomy hack-a-thon event leveraging modern digital engineering tools (analogue to hack-a-sat from cybersecurity domain).

TOPIC 4: EXPLAINABILITY AND HUMAN-MACHINE TEAMING

This topic addresses the effective operations of autonomous systems, and the importance of humans being able to understand the behavior of autonomous systems.

DISCUSSION QUESTIONS:

Trust: Is explainability really needed? Is it a necessary condition for the establishment of trust? Or is it purely for operational reasons (e.g., troubleshooting)?

Content: What if anything needs to be explained to mission control?

Interface: What forms of explanations are needed/most effective?

Costs: Trade-off between performance and explainability

Approach: Is 3rd wave AI (probabilistic logics, neuro-symbolic approaches) going to solve this problem (anytime soon)?

Context: How can explainability support other systems engineering lifecycle processes (not just operations), such as V&V?

TAKEAWAYS:

- Explanations are a useful approach to foster trust in AI agents, which may in turn improve the collaboration and ultimately performance. However, explanations are not a necessary condition for trust. If the user has the opportunity to have many interactions with the agent where the agent shows high accuracy, that can be enough to enable trust. Predictability of the agent's actions also helps foster trust, even if the user does not have a deep understanding of the agent's inner workings. Finally, if the user has been trained to understand how the agent works, explanations may not be needed. But how do we determine which of these approaches is better for a given agent and task?
- Explanations should be tailored to individual needs to maximize their usefulness. Different users have different preferences about the frequency, content and format of the explanations, and learning and adapting to those preferences would significantly improve the interaction. A challenge related to that is that a large dataset (extensive interactions with a large and diverse group of users) would be needed to achieve this goal, and that may not be feasible in the space industry as much as it is with commercial AI agents

targeting routine tasks. How do we solve this problem? Is there something we can learn from the self-driving car industry?

- Future missions will require high levels of autonomy, but at the end of the day there is always a human in or on the loop that needs to be informed of the situation. Different tasks may require different levels of automation, and it seems that there is a trade-off between this level of automation and the amount of control the human has, as well as the invasiveness of the approach (i.e., the amount of data the agent has access to). How do we break the trade-off between automation and control? How do we even measure things like the invasiveness of an agent?

CANDIDATE ACCELERATING ACTIONS:

- Workforce Training. Users of AI agents such as astronauts and mission operators should receive training, including general training in AI, machine learning, and data mining, and specific training about the architecture of the agents they are going to be using. This recommendation has an obvious cost but we believe it to be essential to enable human-machine collaborations and ultimately successful missions.
- Crowdsourcing. Modern approaches to generate training datasets such as crowdsourcing could be considered to alleviate the problem the space industry has with small datasets (few users, little time and opportunities to interact with agents).
- Standards and Guidelines. Standards and guidelines should be developed that help commercial actors develop AI agents that will successfully improve performance, reduce user cognitive workload, improve situational awareness, and foster trust. Those standards and guidelines should be developed based on results from research funded by space programs but also building upon what we are learning from other industries such as self-driving cars. Joint workshops with those industries could be a good venue to discuss these standards and guidelines.

Proposed Next Steps for the AIAA Spacecraft Autonomy Community

This workshop explored four topics of significant importance to the advancement of autonomous system capabilities in space. The questions and takeaway points from this workshop provide a helpful starting point for continued dialogue on this critical subject. 2023 ASCEND and AIAA SciTech Forum present excellent venues for exploring these and other important topic related to spacecraft autonomy. The leadership team is working with key AIAA technical committees (e.g., from the Information Systems Group) to plan a compelling follow-up to this workshop at 2023 ASCEND, building on the progress documented in this report.

The candidate accelerating actions proposed in this report represent seed ideas for AIAA and the broader space community, in their efforts to map out the next steps in the development and deployment of autonomous system capability, for the benefit of our future space exploration, defense, and commercial activities. Fleshing out these next steps is a primary aim of the AIAA's Task Force on Autonomy, Artificial Intelligence, and Machine Learning, which is being stood up under the umbrella of AIAA's Aerospace R&D Domain. The charter, scope, and specific objectives of this task force is the focus of a design sprint activity in mid-2023. Following the design sprint, the task force will conduct a deep dive on topics raised in the AIAA Space Autonomy Summit[1] (and earlier aviation-oriented summit), as well as our autonomy-focused workshops at ASCEND, to formulate technical and policy recommendations.

Finally, this report will be distributed to the AIAA technical committees that are stakeholders in the advancement of spacecraft autonomy, with the intent to stimulate additional activities within those communities and collaborations between them.

Acknowledgments

AIAA would like to thank the following people who participated in the panel and roundtable:

WORKSHOP CHAIRS:

Michel Ingham, NASA Jet Propulsion Laboratory

John Day, Blue Origin

WORKSHOP FACILITATORS AND CONTRIBUTORS:

Richard Doyle, NASA Jet Propulsion Laboratory, retired

Benjamin Seibert, U.S. Space Force

Aaron Celaya, U.S. Space Force

Daniel Selva, Texas A&M University

Timothy Cichan, Lockheed Martin

1 Another opportunity to further the conversation was provided by the AIAA Space Autonomy Summit in Washington, DC, on 16 November 2022. This one-day event brought together national-level thought leaders and stakeholders for space autonomy to initiate discussion on and foster the development of a coherent national strategy for space autonomy.